arXiv:1603.04109v1 [cs.CG] 14 Mar 2016

Combinatorial rigidity of Incidence systems and Application to Dictionary learning*

Meera Sitharam Mohamad Tarifi Menghan Wang

March 15, 2016

Abstract

Given a hypergraph H with m hyperedges and a set Q of m *pinning subspaces*, i.e. globally fixed subspaces in Euclidean space $\mathbb{R}^d$, a *pinned subspace-incidence system* is the pair (H, Q), with the constraint that each pinning subspace in Q is contained in the subspace spanned by the point realizations in $\mathbb{R}^d$ of vertices of the corresponding hyperedge of H. This paper provides a combinatorial characterization of pinned subspace-incidence systems that are *minimally rigid*, i.e. those systems that are guaranteed to generically yield a locally unique realization.

Pinned subspace-incidence systems have applications in the *Dictionary Learning (aka sparse coding)* problem, i.e. the problem of obtaining a sparse representation of a given set of data vectors by learning *dictionary vectors* upon which the data vectors can be written as sparse linear combinations. Viewing the dictionary vectors from a geometry perspective as the spanning set of a subspace arrangement, the result gives a tight bound on the number of dictionary vectors for sufficiently randomly chosen data vectors, and gives a way of constructing a dictionary that meets the bound. For less stringent restrictions on data, but a natural modification of the dictionary learning problem, a further dictionary learning algorithm is provided. Although there are recent rigidity based approaches for low rank matrix completion, we are unaware of prior application of combinatorial rigidity techniques in the setting of Dictionary Learning. We also provide a systematic classification of problems related to dictionary learning together with various algorithms, their assumptions and performance.

1 Introduction

A *pinned subspace-incidence system* (H, Q) is an incidence constraint system specified as a hypergraph H together with a set Q of *pinning subspaces* in $\mathbb{R}^d$, each specified by a collection of basis vectors called *pins*. The pinning subspaces are in one-to-one

*This research was supported in part by the research grant NSF CCF-1117695 and a research gift from SolidWorks. Portions of this paper have appeared on Proceedings of 2014 Canadian Conference on Computational Geometry and Post-Proceedings of 2014 International Workshop on Automated Deduction in Geometry.

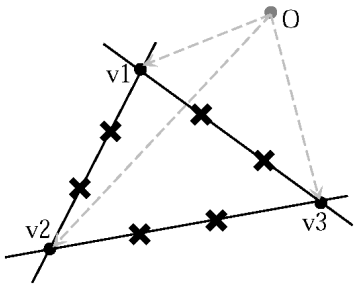

Figure 1: A pinned subspace-incidence framework with $d = 3$, projectivized in $\mathbb{P}^2(\mathbb{R})$.

correspondence with the hyperedges of H. A realization of (H, Q) is a subspace arrangement that assigns vectors in $\mathbb{R}^d$ to the vertices of H. Each hyperedge of H is assigned the subspace spanned by its vertex vectors and contains the associated pinning subspace in Q.

Pinned subspace-incidence systems naturally arise in finding bounds for the *dictionary learning (aka sparse coding)* problem [35] for highly general data, i.e. the problem of obtaining a sparse representation of data vectors by learning *dictionary vectors* upon which the data vectors can be written as sparse linear combinations. This set of dictionary vectors can be viewed from a geometry perspective as the spanning set of a subspace arrangement, where the subspaces contain the data vectors. The subset of data vectors on a given subspace spans a pinning subspace. Thus the solution of dictionary learning problem corresponds to a pinned subspace-incidence system.

The special case of 2-dimensional pinned line-incidence systems was also used in modeling microfibirils in biomaterials such as cellulose and collagen [7]. Each such microfibril is attached to some fixed larger organelle/membrane at one site and crosslinked at two sites with other fibrils, where the *cross-linking* is like an incidence constraint that the crosslinked fibrils can slide against each other while remaining incident. Consequently, they can be modeled using a pinned line-incidence system with H being a graph, where each fibril is modeled as an edge of H with the two cross-linkings as its two vertices, and the attachment is modeled as the corresponding pin.

Previous works on related types of frameworks include pin-collinear body-pin frameworks [23], direction networks [50], slider-pinning rigidity [43], body-cad constraint system [17], k-frames [47, 48], and affine rigidity [15]. All of which involve some form of incidence constraints. However, we are not aware of any previous results on systems that are similar to pinned subspace-incidence systems.

2 Contributions

In this paper, we follow the combinatorial rigidity approaches [6, 47] to give a complete combinatorial characterization for pinned subspace-incidence systems. Specifically,

- We formulate the pinned subspace-incidence systems as a nonlinear algebraic

system $(H, Q)(p)$ and apply classic method of Asimow and Roth [6] by linearizing $(H, Q)(p)$.

- We then apply another well-known method of White and Whiteley [47] to combinatorially characterize the rigidity of the underlying hypergraph H, using the Laplace decomposition of the *rigidity matrix*, which corresponds to a *map-decomposition* [42] of the underlying hypergraph. The polynomial resulting from the Laplace decomposition is called the *pure condition*, which characterizes the conditions that the framework has to avoid for the combinatorial characterization to hold.

To our best knowledge, the only known results with a similar flavor are [17, 30] which characterize the rigidity of Body-and-Cad frameworks. However, these results are dedicated to specific frameworks in 3D instead of arbitrary dimension subspace arrangements and hypergraphs, and their formulation process start directly with the linearized Jacobian (omitting the first bullet alone).

We then apply the combinatorial rigidity result to dictionary learning problems.

- As a corollary of the main result, we give a tight bound on the number of dictionary vectors for sufficiently randomly chosen data vectors, and give a way of constructing a dictionary that meets the bound (see Corollary 14 and 15).

- On the other hand, more common types of data can be handled using a standard preprocessing step such as generalized PCA [46] that converts the dictionary learning problem to a so-called *fitted dictionary learning problem* that yields a specific pinned subspace incidence system where a realization yields a dictionary, followed by recursive decomposition of the underlying pinned subspace-incidence system (referred to as *DR-planning*) to obtain the realization.

- We also provide a systematic classification of problems related to Dictionary Learning together with various approaches, assumptions required and performance (see Section 6.1).

There are some recent applications of rigidity in machine learning [24, 27, 38, 31], specifically for low rank matrix completability. These use the graph of entries of distance matrices and gram matrices, i.e., rigidity with respect to distance and inner product constraints. We are however unaware of applications of rigidity with respect to incidence constraints, or to dictionary learning.

3 Preliminaries

In this section, we introduce the formal definition of pinned subspace-incidence systems and basic concepts in combinatorial rigidity.

A *hypergraph* $H = (V, E)$ is a set V of vertices and a set E of hyperedges, where each hyperedge is a subset of V. The *rank* $r(H)$ of a hypergraph H is the maximum cardinality of any edge in E, i.e. $r(H) = \max_{e_k \in E} s(e_k)$, where $s(e_k)$ denotes the cardinality of the hyperedge e_k. A hypergraph is *s-uniform* if all edges in E have the same cardinality s, where 2-uniform hypergraph is called a *graph* G. A *configuration*

or *realization* of a hypergraph $H = (V, E)$ in $\mathbb{R}^d$ is a mapping from the vertices of H to the vectors in $\mathbb{R}^d$, i.e. $p : V \to \mathbb{R}^d$. When there is no ambiguity, we simply use p_i to denote the vector $p(v_i)$, $p(e_k)$ to denote the set of vectors $\{p(v_i)|v_i \in e_k\}$, and s_k to denote the cardinality $s(e_k)$.

In the following, we use $\langle P \rangle$ to denote the subspace *spanned* by a set P of points in $\mathbb{R}^d$.

Definition 1 (Pinned Subspace-Incidence System)**.** *A* pinned subspace-incidence system *in $\mathbb{R}^d$ is a pair (H, Q), where $H = (V, E, m)$ is a weighted hypergraph of rank $r(H) < d$, and $Q = \{q_1, q_2, \ldots, q_{|E|}\}$ is a set of* pinning subspaces *(subspaces of $\mathbb{R}^d$) in one-to-one correspondence with the hyperedges of H. Here the weight assignment is a function $m : E \to \mathbb{Z}^+$, where $m(e_k)$ denotes the dimension of the pinning subspace q_k associated with the hyperedge e_k. We may write $m(q_k)$ or simply m_k in substitute of $m(e_k)$. Often we ignore the weight m and just refer to the hypergraph (V, E) as H.*

A pinned subspace-incidence framework realizing *the pinned subspace-incidence system (H, Q) is a triple (H, Q, p), where p is a realization of H, such that for all pinning subspaces $q_k \in Q$, q_k is contained in $\langle p(e_k) \rangle$, the subspace spanned by the set of vectors realizing the vertices of the hyperedge e_i corresponding to q_k.*

Since we only care about incidence relations, we projectivize the Euclidean space $\mathbb{R}^d$ to treat the pinned subspace-incidence system in the real projective space $\mathbb{P}^{d-1}(\mathbb{R})$, and refer to the vectors p_i realizing the vertices of H as points in $\mathbb{P}^{d-1}(\mathbb{R})$ using the same notation when the meaning is clear from the context. Figure 1 gives an example of a pinned subspace-incidence framework in the real projective space $\mathbb{P}^{d-1}(\mathbb{R})$ with $d = 3$, where the crosses denote the projectivized pins. As each pinning subspace is spanned by two pins, $m(e) = 2$ for any hyperedge e.

Note: as the pinning subspaces in Q are treated as globally fixed, the *trivial motion group* of a pinned subspace-incidence system (H, Q) reduces to the identity.

Definition 2. *A pinned subspace-incidence system (H, Q) is* independent *if none of the polynomial constraints is in the real ideal generated by others, implying existence of a realization. It is* rigid *if there exist at most finitely many realizations. It is* minimally rigid *if it is both rigid and independent. It is* globally rigid *if there exists at most one realization.*

4 Algebraic Representation and Linearization

In the following, we use $A[R, C]$ to denote a submatrix of a matrix A, where R and C are respectively index sets of the rows and columns contained in the submatrix. In addition, $A[R, \cdot\,]$ represents the submatrix containing row set R and all columns, and $A[\,\cdot\,, C]$ represents the submatrix containing column set C and all rows.

4.1 Representation Using Polynomials

For any hyperedge $e_k = \{v_1^k, v_2^k, \ldots, v_{|e_k|}^k\}$, the subspace $\langle p(e_k) \rangle$ spanned by the point set $\{p_1^k, p_2^k, \ldots, p_{|e_k|}^k\}$ is constrained to contain the pinning subspace q_k associated with

e_k. Recall that q_k is a subspace of dimension $m_k - 1$ in $\mathbb{P}^{d-1}(\mathbb{R})$ spanned by a set of m_k pins $\{x_1^k, x_2^k, \ldots, x_{m_k}^k\}$, and the incidence constraint is equivalent to requiring $\langle p(e) \rangle$ to contain each pin x_l^k, for $1 \leq l \leq m_k$.

Using homogeneous coordinates $p_i^k = [\ p_{i,1}^k \quad p_{i,2}^k \quad \ldots \quad p_{i,d-1}^k\ ]$ and $x_l^k = [\ x_{l,1}^k \quad x_{l,2}^k \quad \ldots \quad x_{l,d-1}^k\ ]$, we write this incidence constraint for each pin x_l^k by letting all the $|e_k| \times |e_k|$ minors of the $|e_k| \times (d-1)$ matrix

$$E_l^k = \begin{bmatrix} p_1^k - x_l^k \\ p_2^k - x_l^k \\ \vdots \\ p_{|e_k|}^k - x_l^k \end{bmatrix}$$

be zero. There are $\binom{d-1}{|e_k|}$ minors, giving $\binom{d-1}{|e_k|}$ equations. Note that any $d - |e_k|$ of these $\binom{d-1}{|e_k|}$ equations are independent and span the rest. So we can write the incidence constraint as $(d - |e_k|)$ independent equations:

$$\det\left(E_l^k[\,\cdot\,, C(t)]\right) = 0, \qquad 1 \leq t \leq d - |e_k| \tag{1}$$

where $C(t)$ denote the following index sets of columns in E:

$$C(t) = \{1, 2, \ldots |e_k| - 1\} \cup \{|e_k| - 1 + t\}, \qquad 1 \leq t \leq d - |e_k|$$

In other words, $C(t)$ contains the first $|e_k|-1$ columns together with column $|e_k|-1+t$.

Now the incidence constraint for the pinning subspace q_k is represented as $m_k(d - |e_k|)$ equations for all the m_k pins $\{x_1^k, x_2^k, \ldots x_{m_k}^k\}$. Consequently, the pinned subspace-incidence problem reduces to solving a system of $\sum_{k=1}^{|E|} m_k(d - |e_k|)$ equations, each of form (1). We denote this algebraic system by $(H, Q)(p) = 0$.

4.2 Linearization and Genericity

We are interested in characterizing *minimal rigidity* of pinned subspace-incidence systems. However, checking independence relative to the ideal generated by the variety is computationally hard and best known algorithms, such as computing Gröbner basis, are exponential in time and space [33]. So we define the notion of rigidity for frameworks (H, Q, p), which is equivalent to maximal rank of the Jacobian of $(H, Q)(p)$.

Definition 3. *A pinned subspace-incidence framework (H, Q, p) is* rigid *if there exists a neighborhood $N(p)$ such that (H, Q, p) is the only framework realizing (H, Q) in $N(p)$. A rigid pinned subspace-incidence framework (H, Q, p) is* minimally rigid *if it is no longer rigid after removing any pin.*

A *generic* framework with respect to a property $\mathcal{P}$, when viewed as a point in an appropriate real space, avoids a measure-zero set $\mathcal{N}_\mathcal{P}$ of the ambient space of frameworks that depends only on the underlying weighted hypergraph. This implies the following notion of genericity:

Definition 4. *A pinned subspace-incidence framework (H, Q, p) is* generic *w.r.t. a property $\mathcal{P}$ if and only if there exists a neighborhood $N(Q, p)$ such that for all frameworks (H, Q', p') with $(Q', p') \in N(Q, p)$, (H, Q', p') satisfies $\mathcal{P}$ if and only if (H, Q, p) satisfies $\mathcal{P}$.*

Furthermore we can define when a property is *generic*, i.e. becomes a property of the hypergraph underlying a geometric constraint system.

Definition 5. *A property $\mathcal{P}$ is* generic *(i.e, becomes a property of the underlying weighted hypergraph alone) if for any weighted hypergraph $H = (V, E, m)$, either all generic (w.r.t. $\mathcal{P}$) frameworks (H, Q, p) satisfy $\mathcal{P}$, or all generic (w.r.t. $\mathcal{P}$) frameworks (H, Q, p) do not satisfy $\mathcal{P}$.*

The primary activity of the area of combinatorial rigidity is to give purely combinatorial characterizations of generic properties $\mathcal{P}$. In practice, the set $\mathcal{N}_\mathcal{P}$ defining genericity is usually not specified, as long as it is of measure zero in the ambient space of frameworks. The measure-zero set $\mathcal{N}_\mathcal{P}$ may include zero-sets of some polynomials called *pure conditions* that appear in the process of drawing such combinatorial characterizations (we will see this in the proof of Theorem 12).

We define the generic rigidity of a pinned subspace-incidence system to be the rigidity of a generic framework.

Definition 6. *A pinned subspace-incidence system (H, Q) is* generically (minimally) rigid *if some generic framework (H, Q, p) realizing (H, Q) is (minimally) rigid.*

4.2.1 Linearization as Rigidity Matrix

Next we follow the approach taken by traditional combinatorial rigidity theory [6, 16] to show that rigidity and independence (based on nonlinear polynomials) of pinned subspace-incidence systems are generically properties of the underlying weighted hypergraph H, and can furthermore be captured by linear conditions in an infinitesimal setting. Specifically, we give a lemma showing that generic rigidity of a pinned subspace-incidence system is equivalent to existence of a full rank *rigidity matrix*, obtained by taking the Jacobian of the algebraic system $(H, Q)(p) = 0$ at a generic framework (H, Q, p) realizing (H, Q).

A *rigidity matrix* of a framework (H, Q, p) is the whose kernel is the infinitesimal motions (flexes) of (H, Q, p). A framework is *infinitesimally rigid* if the rigidity matrix has full rank. To define a rigidity matrix for a pinned subspace-incidence framework (H, Q, p), we take the Jacobian of the algebraic system $(H, Q)(p) = 0$ by taking partial derivatives with respect to the coordinates of p_i's. In the Jacobian, each vertex v_i has $d - 1$ corresponding columns, and each pinning subspace q_k associated with the hyperedge $e_k = \{v_1^k, v_2^k, \ldots, v_{|e_k|}^k\}$ has $m_k(d - |e_k|)$ corresponding rows, where each equation $\det\left(E_l^k[\,\cdot\,, C(t)]\right) = 0$ (1), i.e. equation t of the pin x_l^k, gives the following row (the columns corresponding to vertices not in e_k are all zero):

$$\Bigg[0,\dots,0,\frac{\partial \det E_l^k[\,\cdot\,,C(t)]}{\partial p_{1,1}^k},\frac{\partial \det E_l^k[\,\cdot\,,C(t)]}{\partial p_{1,2}^k},\dots,\frac{\partial \det E_l^k[\,\cdot\,,C(t)]}{\partial p_{1,d-1}^k},0,\dots$$
$$\dots,0,\frac{\partial \det E_l^k[\,\cdot\,,C(t)]}{\partial p_{2,1}^k},\frac{\partial \det E_l^k[\,\cdot\,,C(t)]}{\partial p_{2,2}^k},\dots,\frac{\partial \det E_l^k[\,\cdot\,,C(t)]}{\partial p_{2,d-1}^k},0,\dots$$
$$\dots\dots$$
$$\dots,0,\frac{\partial \det E_l^k[\,\cdot\,,C(t)]}{\partial p_{|e_k|,1}^k},\frac{\partial \det E_l^k[\,\cdot\,,C(t)]}{\partial p_{|e_k|,2}^k},\dots,\frac{\partial \det E_l^k[\,\cdot\,,C(t)]}{\partial p_{|e_k|,d-1}^k},0,\dots,0\Bigg] \quad (2)$$

Let V^k be the matrix whose rows are coordinates of $p_1^k, p_2^k, \dots, p_{|e_k|}^k$:

$$\begin{bmatrix} p_{1,1} & p_{1,2} & \cdots & p_{1,d-1} \\ p_{2,1} & p_{2,2} & \cdots & p_{2,d-1} \\ \vdots & \vdots & \ddots & \vdots \\ p_{|e_k|,1} & p_{|e_k|,2} & \cdots & p_{|e_k|,d-1} \end{bmatrix}$$

Let V_t^k be the $V^k[\,\cdot\,,C(t)]$, i.e. the $|e_k|\times|e_k|$ submatrix of V^k containing only columns in $C(t)$. Let $V_{t,j}^k$ ($1 \le j \le d-1$) be the matrices obtained from V_t^k by replacing the column corresponding to coordinate j with the all-ones vector $(1,1,\dots,1)$ for $j \in C(t)$, and the zero matrix for $j \notin C(t)$. Let $D_{t,j}^k$ be the determinant of $V_{t,j}^k$. Let $b_i^{k,l}$ ($1 \le i \le |e_k|$) be the barycentric coordinates of the pin x_l with respect to the points p_i^k, i.e. $x_l = \sum_{i=1}^{|e_k|} b_i^{k,l} p_i^k$. Now (2) can be rewritten in the following simplified form:

$$r_{t,l}^k = \Big[\, 0,\dots,0,\,0,\, D_{t,1}^k b_1^{k,l},\, D_{t,2}^k b_1^{k,l},\,\dots,\, D_{t,d-1}^k b_1^{k,l},\, 0,\, 0,$$
$$\dots,\, 0,\, 0,\, D_{t,1}^k b_2^{k,l},\, D_{t,2}^k b_2^{k,l},\,\dots,\, D_{t,d-1}^k b_2^{k,l},\, 0,\, 0,\,\dots$$
$$\dots\dots$$
$$\dots,\, 0,\, 0,\, D_{t,1}^k b_{|e_k|}^{k,l},\, D_{t,2}^k b_{|e_k|}^{k,l},\,\dots,\, D_{t,d-1}^k b_{|e_k|}^{k,l},\, 0,\, 0,\,\dots,\, 0\,\Big] \quad (3)$$

Each vertex v_i^k has the entries $D_{t,j}^k b_i^{k,l}, 1 \le j \le d-1$ in its $d-1$ columns, among which exactly $|e_k|$ entries with $j \in C(t)$, i.e. the first $|e_k|-1$ columns together with column $|e_k|-1+t$, are generically non-zero. Note that the terms $D_{t,|e_k|-1+t}^k$ are equal for all t, so we may just use D^k to denote it.

For each $1 \le t \le d-|e_k|$, there are m_k rows as (3), where each pin x_l^k corresponds to the row $r_{t,l}^k$ for $1 \le l \le m_k$. These m_k rows have exactly the same row pattern

except for different $b_i^{k,l}$'s:

$$\begin{array}{c} \\ \\ \\ \\ \\ \end{array}\!\!\begin{array}{c} \begin{array}{cccccccccccc} v_{1,1} & & v_{1,2} & \dots & v_{1,d-1} & & & v_{|e_k|,1} & & v_{|e_k|,2} & \dots & v_{|e_k|,d-1} \end{array} \\ \begin{bmatrix} \dots & D^k_{t,1}b^{k,1}_1 & D^k_{t,2}b^{k,1}_1 & \dots & D^k_{t,d-1}b^{k,1}_1 & \dots & \dots & D^k_{t,1}b^{k,1}_{|e_k|} & D^k_{t,2}b^{k,1}_{|e_k|} & \dots & D^k_{t,d-1}b^{k,1}_{|e_k|} & \dots \\ \dots & D^k_{t,1}b^{k,2}_1 & D^k_{t,2}b^{k,2}_1 & \dots & D^k_{t,d-1}b^{k,2}_1 & \dots & \dots & D^k_{t,1}b^{k,2}_{|e_k|} & D^k_{t,2}b^{k,2}_{|e_k|} & \dots & D^k_{t,d-1}b^{k,2}_{|e_k|} & \dots \\ & & & & & & \ddots & & & & & \\ \dots & D^k_{t,1}b^{k,m_k}_1 & D^k_{t,2}b^{k,m_k}_1 & \dots & D^k_{t,d-1}b^{k,m_k}_1 & \dots & \dots & D^k_{t,1}b^{k,m_k}_{|e_k|} & D^k_{t,2}b^{k,m_k}_{|e_k|} & \dots & D^k_{t,d-1}b^{k,m_k}_{|e_k|} & \dots \end{bmatrix} \end{array}$$

Example 1. *For $d = 4$, consider a pinning subspace q with $m(q) = 2$ associated with the hyperedge $e = \{v_1, v_2\}$. The pinning subspace has the following $m(q)\cdot(d-|e|) = 4$ rows in the simplified Jacobian (the index k is omitted):*

$$\begin{array}{c} \\ t=1,l=1 \\ t=1,l=2 \\ t=2,l=1 \\ t=2,l=2 \end{array}\!\!\begin{array}{c} \begin{array}{ccccccccccc} & v_{1,1} & v_{1,2} & v_{1,3} & & v_{2,1} & v_{2,2} & v_{2,3} & \end{array} \\ \begin{bmatrix} \dots & D_{1,1}b^1_1 & Db^1_1 & 0 & \dots & D_{1,1}b^1_2 & Db^1_2 & 0 & \dots \\ \dots & D_{1,1}b^2_1 & Db^2_1 & 0 & \dots & D_{1,1}b^2_2 & Db^2_2 & 0 & \dots \\ \dots & D_{2,1}b^1_1 & 0 & Db^1_1 & \dots & D_{2,1}b^1_2 & 0 & Db^1_2 & \dots \\ \dots & D_{2,1}b^2_1 & 0 & Db^2_1 & \dots & D_{2,1}b^2_2 & 0 & Db^2_2 & \dots \end{bmatrix} \end{array}$$

We define the *rigidity matrix* $M(H, Q, p)$ or simply $M(p)$ for a pinned subspace-incidence framework (H, Q, p) to be the simplified Jacobian matrix obtained above, where each row has form (3). It is a matrix of size $\sum_k m_k(d - |e_k|)$ by $n(d - 1)$. In general, we use M to denote the rigidity matrix of a generic framework (H, Q, p) with respect to infinitesimal rigidity.

Lemma 7. *Infinitesimal rigidity of a generic subspace-incidence framework (H, Q, p) is equivalent to rigidity of (H, Q, p), thus generic rigidity of the system (H, Q).*

Proof. The proof of Lemma 7 follows the approach taken by traditional combinatorial rigidity [6].

First we show that if a framework (H, Q, p) is generic, infinitesimal rigidity implies rigidity. Consider the polynomial system $(H, Q)(p)$ of equations. The Implicit Function Theorem states that there exists a function g, such that $p = g(Q)$ on some open interval, if and only if the rigidity matrix M has full rank. Therefore, if the framework is infinitesimally rigid, the solutions to the algebraic system are isolated points (otherwise g could not be explicit). Since the algebraic system contains finitely many components, there are only finitely many such solution and each solution is a 0 dimensional point. This implies that the total number of solutions is finite, which is the definition of rigidity.

To show that generic rigidity implies generic infinitesimal rigidity, we take the contrapositive: if a generic framework is not infinitesimally rigid, we show that there is a finite flex. If (H, Q, p) is not infinitesimally rigid, then the rank r of the rigidity matrix M is less than $(d - 1)|V|$. Let E^* be a set of edges in H such that $|E^*| = r$ and the corresponding rows in M are all independent. In $M[E^*, \cdot\,]$, we can find r independent columns. Let p^* be the components of p corresponding to those r independent columns

and $p^{*\perp}$ be the remaining components. The r-by-r submatrix $M[E^*, p^*]$, made up of the corresponding independent rows and columns, is invertible. Then, by the Implicit Function Theorem, in a neighborhood of p there exists a continuous and differentiable function g such that $p^* = g(p^{*\perp})$. This identifies p', whose components are p^* and the level set of g corresponding to p^*, such that $(H, Q)(p') = 0$. The level set defines the finite flexing of the framework. Therefore the system is not rigid. □

5 Combinatorial Rigidity Characterization

5.1 Required Hypergraph Properties

This section introduces pure hypergraph properties and definitions that will be used in stating and proving our main theorem.

Definition 8. *A hypergraph $H = (V, E)$ is $(k, 0)$-sparse if for any $V' \subset V$, the induced subgraph $H' = (V', E')$ satisfies $|E'| \leq k|V'|$. A hypergraph H is $(k, 0)$-tight if H is $(k, 0)$-sparse and $|E| = k|V|$.*

This is a special case of the (k, l)-sparsity condition that was widely studied in the geometric constraint solving and combinatorial rigidity literature. A relevant concept from graph matroids is *map-graph*, defined as follows.

Definition 9. *An* orientation *of a hypergraph is given by identifying as the* tail *of each edge one of its endpoints. The* out-degree *of a vertex is the number of edges which identify it as the tail and connect v to $V - v$. A* map-graph *is a hypergraph that admits an orientation such that the out degree of every vertex is exactly one.*

The following lemma from [42] follows Tutte-Nash Williams [34, 44] to give a useful characterization of $(k, 0)$-tight graphs in terms of maps.

Lemma 10. *A hypergraph H is composed of k edge-disjoint map-graphs if and only if H is $(k, 0)$-tight.*

Our characterization of rigidity of a weighted hypergraph H is based on map-decomposition of a *multi-hypergraph* $\widehat{H}$ obtained from H.

Definition 11. *Given a weighted hypergraph $H = (V, E, m)$, the associated* multi-hypergraph *$\widehat{H} = (V, \widehat{E})$ is obtained by replacing each hyperedge e_k in E with a set E^k of $m_k(d - |e_k|)$ copies of* multi-hyperedges*.*

A *labeling* of a multi-hypergraph $\widehat{H}$ gives a one-to-one correspondence between E^k and the set R^k of $m_k(d - |e_k|)$ rows for the hyperedge e_k in the rigidity matrix M, where the multi-hyperedge corresponding to the row $r^k_{t,l}$ is labeled $e^k_{t,l}$.

Note: an alternative representation commonly adopted in geometric constraint solving [4, 13, 20] is to represent H as a bipartite graph $B(H)$, with $d - 1$ copies of each vertex in V as one of its vertex set, and $m_k(d - |e_k|)$ copies of each hyperedge in E as the other vertex set. A combinatorial rigidity characterization can be equivalently stated using either flow based conditions on $B(H)$ or hypergraph sparsity conditions

on $\widehat{H}$ [21, 22, 25]. However, such an equivalence of the two combinatorial properties has no bearing on the proof of equivalence of either combinatorial property to generic rigidity, an algebraic property. Showing that the combinatorial property generically implies the algebraic property is the substance of the proof of the theorem. This is generally called the "Laman direction" and it is in fact where the hardness of every combinatorial characterization of rigidity lies.

5.2 Characterizing Rigidity

In this section, we apply [47] to give combinatorial characterization for minimal rigidity of pinned subspace-incidence systems.

Theorem 12 (main theorem)**.** *A pinned subspace-incidence system with pins being in general position is generically minimally rigid if and only if:*

(1) The underlying weighted hypergraph $H = (V, E, m)$ satisfies $\sum_{k=1}^{|E|} m_k(d - |e_k|) = (d-1)|V|$, and $\sum_{e_k \in E'} m_k(d-|e_k|) \leq (d-1)|V'|$ for every vertex induced subgraph $H' = (V', E')$. In other words, the associated multi-hypergraph $\widehat{H} = (V, \widehat{E})$ has a decomposition into $(d-1)$ maps.

(2) There exists a labeling of $\widehat{H}$ compatible with the map-decomposition (defined later) *such that in each set E^k of multi-hyperedges,*

(2a) two multi-hyperedges $e^k_{t_1,l_1}$ and $e^k_{t_2,l_2}$ with $l_1 = l_2$ are not contained in the same map in the map-decomposition,

(2b) two multi-hyperedges $e^k_{t_1,l_1}$ and $e^k_{t_2,l_2}$ with $t_1 = t_2$ do not have the same vertex as tail in the map-decomposition.

To prove Theorem 12, we apply Laplace expansion to the determinant of the rigidity matrix M, which corresponds to decomposing the $(d-1, 0)$-tight multi-hypergraph $\widehat{H}$ as a union of $d-1$ maps. We then prove $\det(M)$ is not identically zero by showing that the minors corresponding to each map are not identically zero, as long as a certain polynomial called *pure condition* is avoided by the framework.

A Laplace expansion rewrites the determinant of the rigidity matrix M as a sum of products of determinants (brackets) representing each of the coordinates taken separately. In order to see the relationship between the Laplace expansion and the map-decomposition, we first group the columns of M into $d-1$ column groups C_j according to the coordinates, where columns for the first coordinate of each vertex belong to C_1, columns for the second coordinate of each vertex belong to C_2, etc.

Example 2. *For $d = 4$, consider a pinning subspace q with $m(q) = 2$ associated with the hyperedge $e = \{v_1, v_2\}$. The regrouped rigidity matrix has $d - 1 = 3$ column*

groups, where q has the following 4 rows (the index k is omitted):

$$\begin{array}{c} \\ t=1,l=1 \\ t=1,l=2 \\ t=2,l=1 \\ t=2,l=2 \end{array} \begin{array}{c} \begin{array}{cccccccc} & v_{1,1} & & v_{2,1} & & v_{1,2} & & v_{2,2} \end{array} \\ \left[\begin{array}{cccccc|cccc|cccc} \dots & D_{1,1}b_1^1 & \dots & D_{1,1}b_2^1 & \dots & \dots & Db_1^1 & \dots & Db_2^1 & \dots & & & & \\ \dots & D_{1,1}b_1^2 & \dots & D_{1,1}b_2^2 & \dots & \dots & Db_1^2 & \dots & Db_2^2 & \dots & & & & \\ \dots & D_{2,1}b_1^1 & \dots & D_{2,1}b_2^1 & \dots & & & & & & \dots & Db_1^1 & \dots & Db_2^1 \quad \dots \\ \dots & D_{2,1}b_1^2 & \dots & D_{2,1}b_2^2 & \dots & & & & & & \dots & Db_1^2 & \dots & Db_2^2 \quad \dots \end{array}\right] \end{array} \begin{array}{c} \\ \\ \\ \\ \end{array} \begin{array}{cccc} & & v_{1,3} & & v_{2,3} & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \end{array}$$

We have the following observation on the pattern of the regrouped rigidity matrix.

Observation 1. *In the rigidity matrix M with columns grouped into column groups, a hyperedge e_k has $m_k(d-|e_k|)$ rows, each associated with a multi-hyperedge of e_k in $\widehat{H}$. In a column group j where $j \leq |e_k|-1$, each row associated with e_k contains $|e_k|$ nonzero entries at the columns corresponding to vertices of e_k. In a column group j where $j \geq |e_k|$, there are m_k rows $r_{t,l}^k$ with $|e_k|-1+t=j$, each containing $|e_k|$ nonzero entries at the columns corresponding to vertices of e_k, while the remaining rows are all zero.*

A *labeling of $\widehat{H}$ compatible with a given map-decomposition* can be obtained as following. We start from the last column group of M and associate each column group j with a map in the map-decomposition. For each multi-hyperedge of the map that is a copy of the hyperedge e_k, we pick a row $r_{t,j}^k$ that is not all zero in column groups j and label the multi-hyperedge as $e_{t,j}^k$. By the above observation, this is always possible if each map contains at most m_k multi-hyperedges of the same hyperedge e_k, which must be true if there exists any labeling of $\widehat{H}$ satisfying Theorem 12(2a).

In the Laplace expansion

$$\det(M) = \sum_{\sigma} \left(\pm \prod_{j} \det M[R_j^\sigma, C_j] \right) \tag{4}$$

the sum is taken over all partitions σ of the rows into $d-1$ subsets $R_1^\sigma, R_2^\sigma, \dots, R_j^\sigma, \dots, R_{d-1}^\sigma$, each of size $|V|$. In other words, each summation term of (4) contains $|V|$ rows R_j^σ from each column group C_j. Observe that for any submatrix $M[R_j^\sigma, C_j]$, each row has a common coefficient $D_{t,j}^k$, so

$$\det(M[R_j^\sigma, C_j]) = \left(\prod_{r_{t,j}^k \in R_j^\sigma} D_{t,j}^k \right) \det(M'[R_j^\sigma, C_j])$$

where each row of $M'[R_j^\sigma, C_j]$ is either all zero, or of the pattern

$$[0, \dots, 0, b_1^{k,l}, b_2^{k,l}, 0, \dots, b_{|e_k|}^{k,l}, 0, \dots, 0] \tag{5}$$

with non-zero entries only at the $|e_k|$ indices corresponding to $v_i^k \in e_k$.

For a fixed σ, we refer to a submatrix $M[R_j^\sigma, C_j]$ simply as M_j.

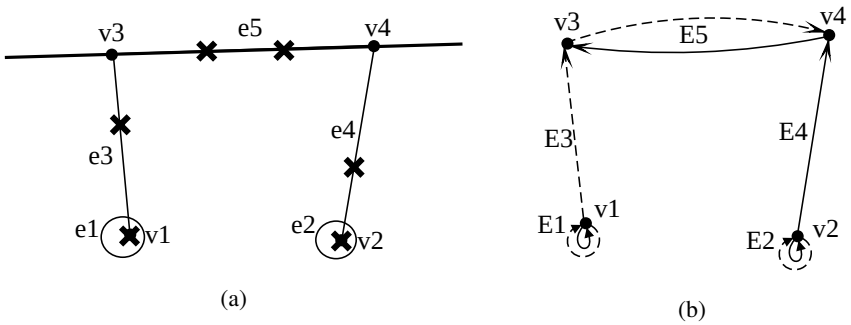

Figure 2: (a) A minimally rigid pinned subspace-incidence system in $d = 3$. (b) A map-decomposition of the multi-hypergraph of the system in (a), where differently drawn multi-hyperedges are in different maps, and the tail vertex of each multi-hyperedge is pointed to by an arrow.

Example 3. *Figure 2a shows a pinned subspace-incidence system in $d = 3$ with 4 vertices and 5 hyperedges, where $e_1 = \{v_1\}, e_2 = \{v_2\}, e_3 = \{v_1, v_3\}, e_4 = \{v_2, v_4\}, e_5 = \{v_3, v_4\}$, and $m_k = 1$ for all $1 \leq k \leq 5$ except that $m_5 = 2$. Figure 2b gives a map-decomposition of the multi-hypergraph $\widehat{H}$ of (a). The labeling of multi-hyperedges is given in the regrouped rigidity matrix* (6)*, where the shaded rows inside the column groups constitute the submatrices M_j^σ in the summation term of the Laplace decomposition corresponding to the map-decomposition. The system is generically minimally rigid, as the map-decomposition and labeling of $\widehat{H}$ satisfies the conditions of Theorem 12.*

	$v_{1,1}$	$v_{2,1}$	$v_{3,1}$	$v_{4,1}$	$v_{1,2}$	$v_{2,2}$	$v_{3,2}$	$v_{4,2}$
$e^1_{1,1}$	$\mathbf{D^1b^1}$							
$e^1_{2,1}$					$\mathbf{D^1b^1}$			
$e^2_{1,1}$		$\mathbf{D^2b^2}$						
$e^2_{2,1}$						$\mathbf{D^2b^2}$		
$e^3_{1,1}$	$D_1^3b_1^3$		$\mathbf{D_1^3b_2^3}$		$D^3b_1^3$		$D^3b_2^3$	
$e^4_{1,1}$		$D_1^4b_1^4$		$D_1^4b_2^4$		$D^4b_1^4$		$\mathbf{D^4b_2^4}$
$e^5_{1,1}$			$D_1^5b_1^{5,1}$	$\mathbf{D_1^5b_2^{5,1}}$			$D_2^5b_1^{5,1}$	$D_2^5b_2^{5,1}$
$e^5_{1,2}$			$D_1^5b_1^{5,2}$	$D_1^5b_2^{5,2}$			$\mathbf{D_2^5b_1^{5,2}}$	$D_2^5b_2^{5,2}$

(6)

Proof of main theorem. First we show the only if direction. For a generically minimally rigid pinned subspace-incidence framework, the rigidity matrix M is generically full rank, so there exists at least one summation term σ in (4) where each submatrix M_j is generically full rank. As the submatrices don't have any overlapping rows with

each other, we can perform row elimination on M to obtain a matrix N with the same rank, where all submatrices M_j are simultaneously converted to a *permuted reduced row echelon form* N_j, where each row in N_j has exactly one non-zero entry β_i^j at a unique column i, In other words, all N_j's can be converted simultaneously to reduced row echelon form by multiplying a permutation matrix on the left of N. Now we can obtain a map-decomposition of $\widehat{H}$ by letting each map j contain multi-hyperedges corresponding to rows of the submatrix N_j, and assigning each multi-hyperedge in map j the vertex i corresponding to the non-zero entry β_i^j in the associated row in N_j as tail. In addition, such a map-decomposition must satisfy the conditions of Theorem 12(2):

Condition (2a): assume two multi-hyperedges $e^k_{t_1,l}$ and $e^k_{t_2,l}$ are in the same map j, i.e. the rows corresponding to these two edges are included in the same submatrix M_j. If $j > s-1$, one of these rows must be all-zero in M_j by Observation 1, contradicting the condition that M_j is full rank. If $j \leq s-1$, both of these rows in M_j will be a multiple of the same row vector (5), contradicting the condition that M_j is full rank.

Condition (2b): note that the rows in M corresponding to multi-hyperedges e^k_{t,l_1} and e^k_{t,l_2} have the exactly the same pattern except for different values of b's. If e^k_{t,l_1} has vertex i as tail, after the row elimination, the column containing b_i^{k,l_1} will become the only non-zero entry of column i for N_{j_1}, while the column containing b_i^{k,l_2} will become zero in all column groups, thus i cannot be assigned as tail for e^k_{t,l_2}.

Next we show the if direction, that the conditions of Theorem 12 imply infinitesimal rigidity.

Given labeled multi-hypergraph $\widehat{H}$ with a map-decomposition satisfying the conditions of Theorem 12, we can obtain summation term σ in the Laplace decomposition (4) according to the labeling of $\widehat{H}$, where each submatrix M_j contain all rows corresponding to the map associated with column group j.

First, it is not hard to show that each submatrix M_j is generically full rank [49]. For completeness, we give a short proof as following. According to the definition of a map-graph, the function $\tau : \widehat{E} \to V$ assigning a tail vertex to each multi-hyperedge is a one-to-one correspondence. We perform symbolic row elimination of the matrix M to simultaneously convert each M_j to its permuted reduced row echelon form N_j, where for each row of N_j, all entries are zero except for the entry $\beta^k_{t,l}$ corresponding to the vertex $\tau(e^k_{t,l})$, which is a polynomial in $b_i^{k,l}$'s in the submatrix M_j. Since M_j cannot contain two rows with the same k and l by Condition (2a), the $b_i^{k,l}$'s in different rows of a same map are independent of each other, $\beta^k_{t,l} \neq 0$ under a generic specialization of $b_i^{k,l}$. Since each row of N_j has exactly one nonzero entry and the nonzero entries from different rows are on different columns, the $|V| \times |V|$ matrix N_j is clearly full rank. Thus M_j must also be generically full rank.

We conclude that

$$\det(M) = \sum_{\sigma} \left(\pm \prod_j \left(\Big(\prod_{r^k_{t,j} \in R^\sigma_j} D^k_{t,j} \Big) \det M'[R^\sigma_j, C_j] \right) \right) \tag{7}$$

where the sum is taken over all σ corresponding to a map-decomposition of $\widehat{H}$. Generically, the summation terms of the sum (7) do not cancel with each other, since $\det(M'[R^\sigma_j, C_j])$

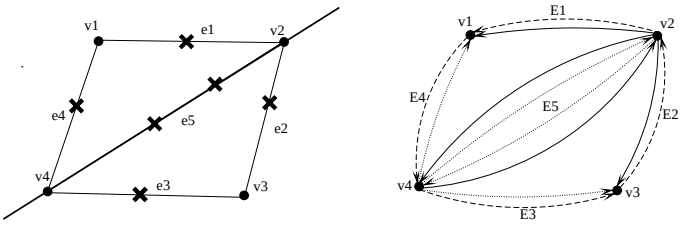

Figure 3: (a) A pinned subspace-incidence system in $d = 4$. (b) A map-decomposition of the multi-hypergraph of the system in (a), where multi-hyperedges with different patterns are in different maps, and the tail vertex of each multi-hyperedge is pointed to by an arrow.

are independent of the multi-linear coefficients $\prod_{r^k_{t,j} \in R^\sigma_j} D^k_{t,j}$ because of the requirement that the pins are in general position, and any two rows of M are independent by Condition (2b). This implies that $\widehat{M}$ is generically full rank.

The polynomial (7) gives the pure condition characterizing non-generic frameworks. □

5.2.1 Pure condition

The pure condition (7) obtained in the proof of Theorem 12 vanishes at a measure-zero subset of frameworks. For those frameworks, the above characterization fails. However, the geometric meaning of the pure condition is not completely clear.

Note that there exist combinatorial types of underlying weighted hypergraphs, for example, H with a subgraph (V', E'), $|V'| < d$ such that $\sum_{e_k \in E'} m_k > |V'|$, that force the pins to lie in a non-generic position no matter how the vertices are realized. We rule out such systems by the requirement that pins being in general position in the statement of Theorem 12.

As an example, Figure 3a shows a pinned subspace-incidence system in $d = 4$ with 4 vertices and 5 hyperedges, where m_k is 2 for $k = 5$ and is 1 otherwise. A map-decomposition of the multi-hypergraph $\widehat{H}$ of the system is given in Figure 3b, and we can easily find a labeling of $\widehat{H}$ satisfying Conditions (1) and (2) in Theorem 12. However, the system is overconstrained, as generically the four pins associated with e_1, e_4 and e_5 will not fall on the same plane in $\mathbb{P}^3(\mathbb{R})$. This situation is both captured by the pure condition and is ruled by the requirement of Theorem 12 that the pins being in general position.

Figure 4 shows a more standard non-generic example captured by the pure condition. Frameworks (a) and (b) are two pinned subspace incidence frameworks with $d = 4$, and they have the same underlying weighted hypergraph satisfying the combinatorial characterization of the main theorem. However, framework (a) is minimally rigid but (b) is not, since the pins on hyperedges e_1 and e_2 in (b) lie on the same line. Evaluating the pure condition at framework (b) shows that it is not generic.

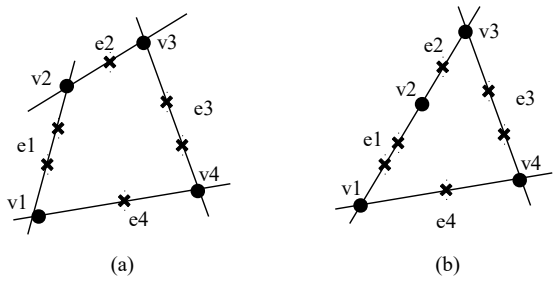

Figure 4: Two pinned subspace-incidence frameworks in $d = 4$ with the same underlying weighted hypergraph, where (a) is generic but (b) is not generic.

6 Application: Dictionary Learning

Dictionary Learning (aka Sparse Coding) is the problem of obtaining a sparse representation of data points, by learning *dictionary vectors* upon which the data points can be written as sparse linear combinations. The Dictionary Learning problem arises in various context(s) such as signal processing and machine learning.

Problem 1 (Dictionary Learning)**.** *A point set $X = [x_1 \dots x_m]$ in $\mathbb{R}^d$ is said to be* s-represented *by a* dictionary $D = [v_1 \dots v_n]$ *for a given* sparsity $s < d$*, if there exists $\Theta = [\theta_1 \dots \theta_m]$ such that $x_i = D\theta_i$, with $\|\theta_i\|_0 \leq s$ (i.e. θ_i has at most s non-zero entries). Given an X known to be s-represented by an unknown* dictionary D *of size $|D| = n$,* Dictionary Learning *is the problem of finding any dictionary $\acute{D}$ satisfying the properties of D, i.e. $|\acute{D}| \leq n$, and there exists $\acute{\theta}_i$ such that $x_i = \acute{D}\acute{\theta}_i$ for all $x_i \in X$.*

The dictionary under consideration is usually *overcomplete*, with $n > d$. However we are interested in asymptotic performance with respect to all four variables n, m, d, s. Typically, $m \gg n \gg d > s$. Both cases when s is large relative to d and when s is small relative to d are interesting.

We understand the Dictionary Learning problem from an intrinsically geometric point of view. Notice that each $x \in X$ lies in an s-dimensional subspace $\mathfrak{S}_D(x)$, which is the span of s vectors $v \in D$ that form the *support* of x. The resulting *s-subspace arrangement* $S_{X,D} = \{\mathfrak{S}_D(x) : x \in X\}$ has an underlying labeled (multi)hypergraph $H(S_{X,D}) = (\mathfrak{I}(D), \mathfrak{I}(S_{X,D}))$, where $\mathfrak{I}(D)$ denotes the index set of the dictionary D and $\mathfrak{I}(S_{X,D})$ is the set of (multi)hyperedges over the indices $\mathfrak{I}(D)$ corresponding to the subspaces $\mathfrak{S}_D(x)$. The word "multi" appears because if $\mathfrak{S}_D(x_1) = \mathfrak{S}_D(x_2)$ for data points $x_1, x_2 \in X$ with $x_1 \neq x_2$, then that subspace is multiply represented in $S_{X,D}$ (resp. $\mathfrak{I}(S_{X,D})$) as $\mathfrak{S}_D(x_1)$ and $\mathfrak{S}_D(x_2)$.

Note that there could be many dictionaries D and for each D, many possible subspace arrangements $S_{X,D}$ that are solutions to the Dictionary Learning problem.

6.1 Systematic Classification of Problems Closely Related to Dictionary Learning and Previous Approaches

A closely related problem to Dictionary Learning is the Vector Selection (aka sparse recovery) problem, which finds a representation of input data in a known dictionary D.

Problem 2 (Vector Selection)**.** *Given a dictionary $D \in \mathbb{R}^{d \times n}$ and an input data point $x \in \mathbb{R}^d$, the* Vector Selection problem *asks for $\theta \in \mathbb{R}^n$ such that $x = D\theta$ with $\|\theta\|_0$ minimized.*

That is, θ is a sparsest support vector that represents x as linear combinations of the columns of D.

An optimization version of Dictionary Learning can be written as:

$$\min_{D \in \mathbb{R}^{d \times n}} \max_{x_i} \min \|\theta_i\|_0 : x_i = D\theta_i.$$

In practice, it is often relaxed to the Lagrangian $\min \sum_{i=0}^{m} (\|x_i - D\theta_i\|_2 + \lambda \|\theta_i\|_1)$.

Several traditional Dictionary Learning algorithms work by *alternating minimization*, i.e. iterating the following two steps [41, 36, 35]:

1. Starting from an initial estimation of D, solving the Vector Selection problem for all data points X to find a corresponding Θ. This can be done using any vector selection algorithm, such as basis pursuit from [10].

2. Given Θ, updating the dictionary estimation by solving the optimization problem is now convex in D. For an overcomplete dictionary, the general Vector Selection problem is ill defined, as there can be multiple solutions for a data point x. Overcoming this by framing the problem as a minimization problem is exceedingly difficult. Indeed under generic assumptions, the Vector Selection problem has been shown to be NP-hard by reduction to the Exact Cover by 3-set problem [37]. One is then tempted to conclude that Dictionary Learning is also NP-hard. However, this cannot be directly deduced in general, since even though adding a witness D turns the problem into an NP-hard problem, it is possible that the Dictionary Learning solves to produce a different dictionary $\acute{D}$.

On the other hand, if D satisfies the condition of being a *frame*, i.e. for all θ such that $\|\theta\|_0 \leq s$, there exists a δ_s such that $(1 - \delta_s) \leq \frac{\|D\theta\|_2^2}{\|\theta\|_2^2} \leq (1 + \delta_s)$, it is guaranteed that the sparsest solution to the Vector Selection problem can be found via L_1 minimization [11, 9].

One popular alternating minimization method is the Method of Optimal Dictionary (MOD) [12], which follows a two step iterative approach using a maximum likelihood formalism, and uses the pseudoinverse to compute D: $D^{(i+1)} = X{\Theta^{(i)}}^T(\Theta^n{\Theta^i}^T)^{-1}$. The MOD can be extended to Maximum A-Posteriori probability setting with different priors to take into account preferences in the recovered dictionary.

Similarly, k-SVD [3] uses a two step iterative process, with a Truncated Singular Value Decomposition to update D. This is done by taking every atom in D and applying SVD to X and Θ restricted to only the columns that have contribution from that atom. When D is restricted to be of the form $D = [B_1, B_2 \dots B_L]$ where B_i's are orthonormal matrices, a more efficient pursuit algorithm is obtained for the sparse coding stage using a block coordinate relaxation.

Though alternating minimization methods work well in practice, there is no theoretical guarantee that the their results will converge to a true dictionary. Several recent works give provable algorithms under stronger constraints on X and D. Spielman et. al [40] give an L_1 minimization based approach which is provable to find the exact dictionary D, but requires D to be a basis. Arora et. al [5] and Agarwal et. al [2] independently give provable non-iterative algorithms for learning approximation of overcomplete dictionaries. Both of their methods are based on an overlapping clustering approach to find data points sharing a dictionary vector, and then estimate the dictionary vectors from the clusters via SVD. The approximate dictionary found using these two algorithms can be in turn used in iterative methods like k-SVD as the initial estimation of dictionary, leading to provable convergence rate [1]. However, these overlapping clustering based methods require the dictionaries to have the *pairwise incoherence* property which is much stronger than the frame property.

By imposing a systematic series of increasingly stringent constraints on the input, we classify previous approaches to Dictionary Learning as well as a whole set of independently interesting problems closely related to Dictionary Learning. A summary of the input conditions and results of these different types of Dictionary Learning approaches can be found in Table 1.

A natural restriction of the general Dictionary Learning problem is the following. We say that a set of data points X *lies on* a set S of s-dimensional subspaces if for all $x_i \in X$, there exists $S_i \in S$ such that $x_i \in S_i$.

Problem 3 (Subspace Arrangement Learning)**.** *Let X be a given set of data points that are known to lie on a set S of s-dimensional subspaces of $\mathbb{R}^d$, where $|S|$ is at most k. (Optionally assume that the subspaces in S have bases such that their union is a frame).* Subspace arrangement learning *finds any subspace arrangement $\acute{S}$ of s-dimensional subspaces of $\mathbb{R}^d$ satisfying these conditions, i.e. $|\acute{S}| \leq k$, X lies on $\acute{S}$, (and optionally the union of the bases of $\acute{S}_i \in \acute{S}$ is a frame).*

There are several known algorithms for learning subspace arrangements. Random Sample Consensus (RANSAC) [45] is an approach to learning subspace arrangements that isolates, one subspace at a time, via random sampling. When dealing with an arrangement of k s-dimensional subspaces, for instance, the method samples $s + 1$ points which is the minimum number of points required to fit an s-dimensional subspace. The procedure then finds and discards inliers by computing the residual to each data point relative to the subspace and selecting the points whose residual is below a certain threshold. The process is iterated until we have k subspaces or all points are fitted. RANSAC is robust to models corrupted with outliers. Another method called Generalized PCA (GPCA) [46] uses techniques from algebraic geometry for subspace clustering, finding a union of k subspaces by factoring a homogeneous polynomial of degree k that is fitted to the points $\{x_1 \dots x_m\}$. Each factor of the polynomail represents the normal vector to a subspace. We note that GPCA can also determine k if it is unknown.

The next problem is obtaining a minimally sized dictionary from a subspace arrangement.

Problem 4 (Smallest Spanning Set for Subspace Arrangement)**.** *Let S be a given set of s-dimensional subspaces of $\mathbb{R}^d$ specified by giving their bases. Assume their intersections are known to be s-represented by a set I of vectors with $|I|$ at most n. Find any set of vectors $\acute{I}$ that satisfies these conditions.*

The smallest spanning set is not necessarily unique in general, and is closely related to the intersection semilattice of subspace arrrangement [8, 14]. Furthermore, under the condition that the subspace arrangement comes from a frame dictionary, the smallest spanning set is the union of: (a) the smallest spanning set I of the pairwise intersection of all the subspaces in S; (b) any points outside the pairwise intersections that, together with I, completely s-span the subspaces in S. This directly leads to a recursive algorithm for the smallest spanning set problem.

The *fitted dictionary learning* problem is the version of dictionary learning where the underlying hypergraph $H(S_{X,D})$ is specified.

Problem 5 (Fitted Dictionary Learning)**.** *Let X be a given set of data points in $\mathbb{R}^d$. For an unknown dictionary $D = [v_1, \ldots, v_n]$ that s-represents X, we are given the hypergraph $H(S_{X,D})$ of the underlying subspace arrangement $S_{X,D}$. Find any dictionary $\acute{D}$ of size $|\acute{D}| \leq n$, that is consistent with the hypergraph $H(S_{X,D})$.*

When X contains sufficiently dense data to solve Problem 3, Dictionary Learning reduces to problem 4, i.e. we can use the following two-step procedure to solve the Dictionary Learning problem:

- Learn a Subspace Arrangement S for X (instance of Problem 4).
- Recover D by either finding the smallest Spanning Set of S (instance of Problem 3), or a fitted dictionary learning (instance of Problem 5).

Note that it is not true that the decomposition strategy should always be applied for the same sparsity s. The decomposition starts out with the minimum given value of s and is reapplied with iteratively higher s if a solution has not be obtained.

A summary of the input conditions and results of these different types of Dictionary Learning problems from this section can be found in Table 1.

6.2 New Bounds and Algorithms for Dictionary Learning for Random Data via Pinned Subspace-Incidence Systems

The following corollary of Theorem 12 gives a tight bound on dictionary size for generic data points, which leads to an algorithm for finding a dictionary provided the bounds hold.

Corollary 13 (Dictionary size tight bound for generic data)**.** *Given a set of m points $X = \{x_1, .., x_m\}$ in $\mathbb{R}^d$, generically there is a dictionary D of size n that s-represents X only if $(d-s)m \leq (d-1)n$. Conversely, if $(d-s)m = (d-1)n$ and the supports of x_i (the nonzero entries of the θ_i's) are known to form a $(d-1, 0)$-tight hypergraph H, then generically, there is at least one and at most finitely many such dictionaries.*

	Traditional Dictionary Learning				Dictionary Learning via subspace arrangement and spanning set	Dictionary Learning for Segmented Data	Fitted Dictionary Learning (this paper)
	Alternating Minimization Approaches	Spielman et. al [40]	Arora et. al [5], Agarwal et. al [2]	For generic data (this paper)			
Input and Conditions	D satisfies frame property	X generated from hidden dictionary D and certain distribution of Θ		Generic data points X	X with promise that each subspace / dictionary support set is shared by sufficiently many data points in X	Partitioned / segmented Data X	X with underlying hypergraph specified
		D is a basis	D is pairwise incoherent				
Minimum m guaranteeing existence of a locally unique dictionary of a given size n	Question 4	$O(n \log n)$	$O(n^2 \log^2 n)$	$\frac{d-s}{d-1}n$ (Corollary 14); Unknown for general position data (Question 2)	Minimum number of points to guarantee a unique subspace arrangement that will give a spanning set of size n	Question 1	$\frac{d-s}{d-1}n$, where the underlying hypergraph satisfy Theorem 12
Dictionary Learning algorithms	MOD, k-SVD, etc.	Algorithm from [40]	Algorithms from [5, 2]	Straightforward algorithm (Corollary 15)	Subspace Arrangement Learning Algorithms (Problem 3) and Spanning Set Finding (Problem 4)	Question 1 and Spanning Set Finding (Problem 4)	Realization using DR-plan (Theorem 17)
Minimum m guaranteeing efficient dictionary learning	Unknown	$O(n^2 \log^2 n)$		Unknown	Unknown	Unknown	Unknown
Illustrative example					(a) 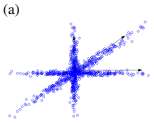	(b) 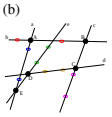	(c) 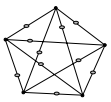

Table 1: Classification of Problems

Quantifying the term "generically" in Corollary 13 yields Corollaries 14 and 15 below.

Corollary 14 (Lower bound for random data)**.** *Given a set of m points $X = \{x_1, .., x_m\}$ picked from a distribution ρ with respect to which nongenericity has measure zero, a dictionary D that s-represents X has size at least $(\frac{d-s}{d-1})m$ with probability 1. In other words, $|D| = \Omega(X)$ if s and d are constants.*

Corollary 15 (Straightforward Dictionary Learning Algorithm)**.** *Given a set of m points $X = [x_1 \dots x_m]$ picked from a distribution ρ with respect to which nongenericity has measure zero, there is a straightforward pebble-game [25, 29, 28] based algorithm constructs a dictionary $D = [v_1 \dots v_n]$ that s-represents X, where $n = \left(\frac{d-s}{d-1}\right) m$. The time complexity of the algorithm is $O(m)$ when we treat d and s as constants.*

The key idea of the algorithm for Corollary 15 is that we can choose a convenient underlying hypergraph $H(S_{X,D})$. Thus the algorithm has two parts: (1) constructing $H(S_{X,D})$ and (2) constructing the s-subspace arrangement $S_{X,D}$ and the dictionary D.

(1) Algorithm for constructing the underlying hypergraph $H(S_{X,D})$ for a hypothetical s-subspace arrangement $S_{X,D}$:

The algorithm works in three stages to construct a expanded mutli-hypergraph $\hat{H}(S_{X,D})$:

1. We start by constructing a minimal minimally rigid hypergraph $H_0 = (V_0, E_0)$, using the *pebble game algorithm* introduced below. Here $|V_0| = k(d-s)$, $|E_0| = k(d-1)$, where k is the smallest positive integer such that $\binom{k(d-s)}{s} \geq k(d-1)$, so it is possible to construct E_0 such that no more than one hyperedge in E_0 containing the same set of vertices in V_0. The values $|V_0|$ and $|E_0|$ are constants when we think of d and s as constants.

2. We use the pebble game algorithm to append a set V_1 of $d-s$ vertices and a set E_1 of $d-1$ hyperedges to H_0, such that each hyperedge in E_1 contains at least one vertex from V_1, and the obtained graph H_1 is still minimally rigid. The subgraph B_1 induced by E_1 has vertex set $V_{B_1} = V_1 \bigcup V_B$, where $V_B \subset V_0$. We call the vertex set V_B the *base vertices* of the construction.

3. Each of the following construction step i appends a set V_i of $d-s$ vertices and a set E_i of $d-1$ hyperedges such that the subgraph B_i induced by E_i has vertex set $V_i \bigcup V_B$, and B_i is isomorphic to B_1. In other words, at each step, we directly append a basic structure the same as (V_1, E_1) to the base vertices V_B. It is not hard to verify that the obtained graph is still minimally rigid.

The *pebble game algorithm* by [42] works on a fixed finite set V of vertices and constructs a (k, l)-sparse hypergraph. Conversely, any (k, l)-sparse hypergraph on vertex set V can be constructed by this algorithm. The algorithm initializes by putting k pebbles on each vertex in V. There are two types of moves:

- *Add-edge*: adds a hyperedge e (vertices in e must contain at least $l+1$ pebbles), removes a pebble from a vertex v in e, and assign v as the tail of e;
- *Pebble-shift*: for a hyperedge e with tail v_2, and a vertex $v_1 \in e$ which containing at least one pebble, moves one pebble from v_1 to v_2, and change the tail of e to v_1.

At the end of the algorithm, if there are exactly l pebbles in the hypergraph, then the hypergraph is (k,l)-tight.

Our algorithm runs a slightly modified pebble game algorithm to find a $(d-1,0)$-tight expanded mutli-hypergraph. We require that each add-edge move adding $(d-s)$ copies of a hyperedge e, so a total of $d-s$ pebbles are removed from vertices in e. Additionally, the multiplicity of a hyperedge, not counting the expanded copies, cannot exceed 1. For constructing the basic structure of Stage 2, the algorithm initializes by putting $d-1$ pebbles on each vertex in V_1. In addition, an add-edge move can only add a hyperedge that contains at least one vertex in V_1, and a pebble-shift move can only shift a pebble inside V_1.

The pebble-game algorithm takes $O\left(s^2|V_0|\binom{|V_0|}{s}\right)$ time in Step 1, and $O\left(s^2\left(|V_0|+(d-s)\right)\binom{|V_0|+(d-s)}{s}\right)$ time in Step 2. Since the entire underlying hypergraph $H(S_{X,D})$ has $m=|X|$ edges, Step 3 will be iterated $O(m/(d-1))$ times, and each iteration takes constant time. Therefore the overall time complexity for constructing $H(S_{X,D})$ is

$$O\left(s^2\left(|V_0|+(d-s)\right)\binom{|V_0|+(d-s)}{s}+(m/(d-1))\right)$$

which is $O(m)$ when d and s are regarded as constants.

(2) Algorithm for constructing the s-subspace arrangement $S_{X,D}$ and the dictionary D:

The construction of the s-subspace arrangement $S_{X,D}$ naturally follows from the construction of the underlying hypergraph $H(S_{X,D})$. For the initial hypergraph H_0, we get a pinned subspace-incidence system $(H_0,X_0)(D_0)$ by arbitrarily choose $|X_0|=|E_0|$ pins from X. Similarly, for Step 2 and each iteration of Step 3, we form a pinned subspace-incidence system $(B_i,X_i)(D_i)$ by arbitrarily choosing $|X_i|=d-1$ pins from X.

Given X_0, we know that the rigidity matrix – of the s-subspace framework $H_0(S_{X_0,D_0})$ – with indeterminates representing the coordinate positions of the points in D_0 – generically has full rank (rows are maximally independent), under the pure conditions of Theorem 12; in which case, the original algebraic subsystem $(H_0,X_0)(D_0)$ (whose Jacobian is the rigidity matrix), with X_0 plugged in, is guaranteed to have a (possibly complex) solution and only finitely many solutions for D_0. Since the pure conditions fail only on a measure-zero subset of the space of pin-sets X_0, where each pin is in S^{d-1}, it follows that if the pins in X_0 are picked uniformly at random from S^{d-1} we know such a solution exists for D_0 (and S_{X_0,D_0}) and can be found by solving the algebraic system $H_0(S_{X_0,D_0})$.

Once we have solved $(H_0, X_0)(D_0)$, for each following construction step i, B_i is also rigid since coordintate positions of the vertices in V_B have been fixed So similarly, we know a solution exists for D_i (and S_{X_i,D_i}) and can be found by solving the algebraic system $B_i(S_{X_i,D_i})$, which is of constant size $O(d)$. Although there can be more than one choice of solution for each step, since every construction step is based on base vertices V_B, the solution of one step will not affect any other steps, so generically any choice will result in a successful solution for the entire construction sequence, and we obtain D by taking the union of all D_i's.

When we regard d and s as constants, the time complexity for Stage (2) is the constant time for solving the size $O(|V_0|)$ algebraic system $(H_0, X_0)(D_0)$, plus $O(m/(d-1))$ times the constant time for solving the size $O(d)$ system $(B_i, X_i)(D_i)$, that is $O(m)$ in total.

Therefore the overall time complexity of the dictionary learning algorithm is $O(m)$.

6.3 Dictionary Learning for Commonly Occurring Data Using DR-Planning

In practical dictionary learning problems, the data set is usually overconstrained, making it impossible to apply the algorithm given above.

As one solution for such dictionary learning problems, recall the two-step procedure described in Section 6.1, where we first use standard preprocessing algorithms such as RANSAC and GPCA to learn a subspace arrangement from the data set (Problem 3), and then constructing a fitted dictionary learning problem (Problem 5) based on the subspace arrangement learnt. Now we can obtain a dictionary by realizing the (possibly overconstrained) pinned subspace incidence-system corresponding to the fitted dictionary learning problem.

To find a realization to a geometric constraint system, the straightforward method is to directly find the real solutions to the entire multivariate polynomial system. However, such an approach requires double exponential time in the number of variables. Thus, it is crucial to use recursive *Decomposition-Recombination (DR-) plans* [19, 18, 26, 39] , which decompose the original system into locally rigid subsystems and recombines solutions of subsystems to get a solution for the original system. For a pinned subspace-incidence system (H, Q), a DR-plan is formally defined for the underlying multi-hypergraph $\hat{H}$ (recall the definition from Section 5.1) as following:

Definition 16 (DR-plan)**.** *A* decomposition-recombination (DR-)plan *of the underlying multi-hypergraph $\hat{H}$ of a pinned subspace-incidence system (H, Q) in $\mathbb{R}^d$ is a forest that has the following properties:*

1. *Each node represents a connected vertex-induced subgraph of $\hat{H}$ that is either rigid, or a single hyperedge.*
2. *A root node represents a maximal subgraph of $\hat{H}$.*
3. *A node represents the subgraph of $\hat{H}$ induced by the union of the node's children.*
4. *A leaf node represents a single hyperedge.*

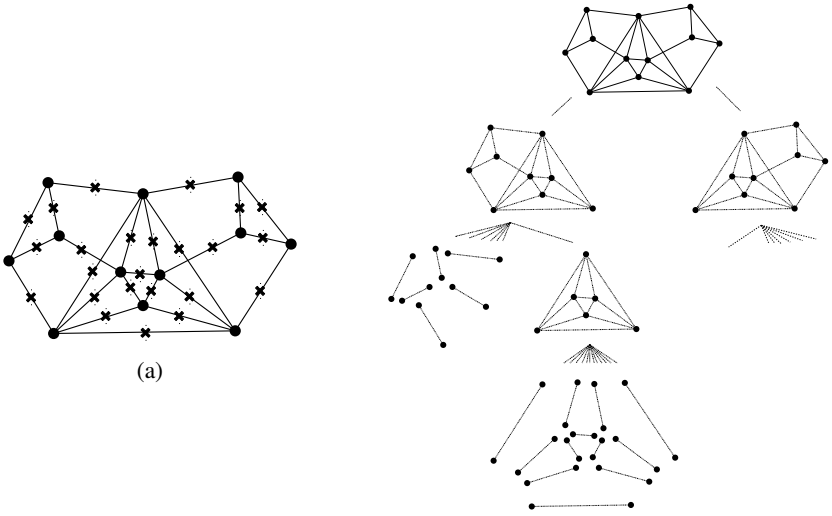

(a) (b)

Figure 5: (a) A generically minimally rigid pinned subspace-incidence system with $d = 3$. (b) The optimal DR-plan of (a) obtained using Theorem 17.

With the use of a DR-plan, the complexity of realizing a geometric constraint system is determined by the maximum *fan-in* of the DR-plan, i.e. the maximum number of children nodes of any node. Finding an *optimal DR-plan* i.e. one that minimizes the maximum fan-in, however, is usually hard for general geometric constraint systems.

In the recent work [7], it is shown that an optimal DR-plan can be efficiently found for a large class of minimally rigid geometric constraint systems with $O(n^3)$ time complexity, leading to efficient realization. Specifically, this result holds for minimally rigid pinned subspace-incidence systems, as stated in the following theorem:

Theorem 17 ([7])**.** *For generically minimally rigid pinned subspace-incidence, there exists a DR-plan of the underlying multi-hypergraph satisfying the following two additional properties, and any such DR-plan is optimal:*

1. *Children are connected, generically rigid, vertex-maximal proper subgraphs of the parent.*
2. *If all pairs of rigid vertex-maximal proper subgraphs intersect trivially then all of them are children, otherwise exactly two that intersect non-trivially are children.*

Figure 5 shows an example of the optimal DR-plan found by applying Theorem 17 to a generically minimally rigid pinned subspace-incidence system with 12 vertices and 24 pins, where $d = 3$. The DR-plan has the optimal maximum fan-in 12.

However, for overconstrained systems, the optimal DR-planning problem was shown to be NP-hard even for bar-joint systems [32], and we strongly conjecture it is also NP-hard for pinned subspace-incidence systems. One possible workaround is to first apply

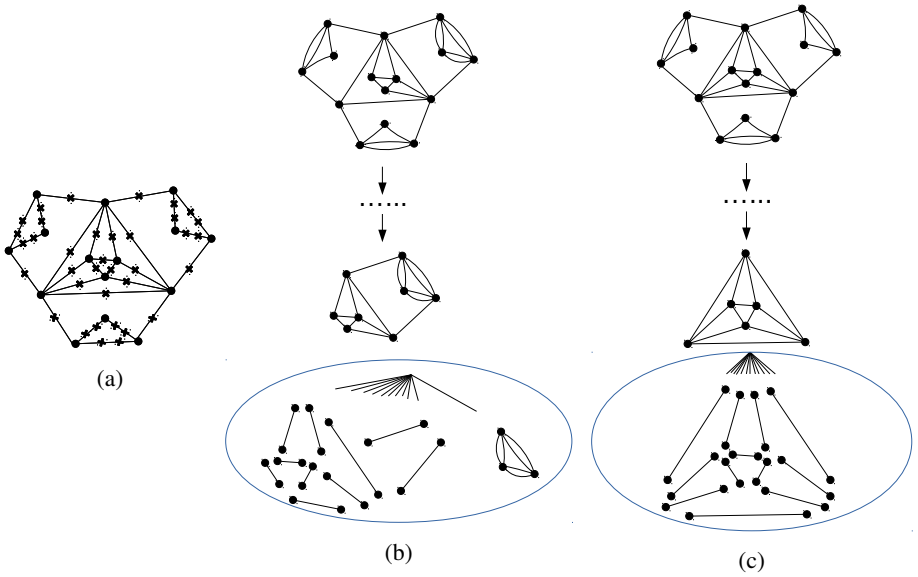

Figure 6: (a) A overconstrained pinned subspace-incidence system with $d = 3$. (b) A maximal minimally rigid subgraph with maximum fan-in 11 (optimal). (c) A maximal minimally rigid subgraph with maximum fan-in 11

the pebble-game algorithm from [25, 29] to an overconstrained system (H, Q) to find a maximal, generically minimally rigid subsystem, and then apply the algorithm from [7] to find an optimal DR-plan for the subsystem. However, by doing this we are forgoing the advantages of smaller sized DR-plan which is usually provided by overconstrained cases. We also note that different maximal, generically minimally rigid subsystems of (H, Q) may still have different maximum fan-ins. As an example, Figure 6(a) shows a overconstrained pinned subspace-incidence system with $d = 3$. Figure 6 (b) and (c) are DR-plans of two maximal, generically minimally rigid subsystems of (a), where (b) has the optimal maximum fan-in 11, while (c) has a maximum fan-in of 12.

7 Open Questions

A natural restriction of the two step problem in Section 6.1, which learns subspace arrangement followed by spanning set is the following, where the data set X is given in *support-equivalence* classes. For a given subspace t in the subspace arrangement $S_{X,D}$ (respectively hyperedge h in the hypergraph's edge-set $\mathfrak{I}(S_{X,D})$), let $X_t = X_h \subseteq X$ be the equivalence class of data points x such that $\mathfrak{S}_D(x) = t$. We call the data points x in a same X_h as *support-equivalent*.

Question 1 (Dictionary Learning for Partitioned Data)**.** *(1) What is the minimum size of X and X_i's (representing data X partitioned into X_i's) guaranteeing that there exists a locally unique dictionary D for a s-subspace arrangement $S_{X,D}$ satisfying $|D| \leq n$,*

and X_i represents the support-equivalence classes of X with respect to D?
(2) How to find such a dictionary D?

With regard to the problem of minimizing $|D|$, very little is known for simple restrictions on X. For example the following question is open.

Question 2. *Given a general position assumption on X, what is the best lower bound on $|D|$ for Dictionary Learning? Conversely, are smaller dictionaries possible than indicated by Corollary 14 (see Section 6.2) under such an assumption?*

Question 2 gives rise to the following pure combinatorics open question closely related to the intersection semilattice of subspace arrrangement [8, 14].

Question 3. *Given weights $w(S) \in \mathbb{N}$ assigned to size-s subsets S of $[n]$. For $T \subseteq [n]$ with $|T| \neq s$,*

$$w(T) = \begin{cases} 0 & |T| < s \\ \sum_{S \subset T, |S|=s} w(S) & |T| > s \end{cases}$$

Assume additionally the following constraint holds: for all subsets T of $[n]$ with $s \leq |T| \leq d$, $w(T) \leq |T| - 1$. Can one give a nontrivial upper bound on $w([n])$?

The combinatorial characterization given by Theorem 12 leads to the following question for general Dictionary Learning.

Question 4. *What is the minimum size of a data set X such that the Dictionary Learning for X has a locally unique solution dictionary D of a given size? What are the geometric characteristics of such an X?*

8 Conclusion

In this paper, we studied the generic rigidity of pinned subspace-incidence systems, and obtained a combinatorial characterization of generic minimal rigidity for pinned subspace-incidence systems. We then showed the application of pinned subspace-incidence systems in Dictionary Learning, giving a tight bound on dictionary size in a specific setting as well as new dictionary learning algorithms in this and other general settings. We provided a systematic classification of problems related to dictionary learning together with various algorithms, their assumptions and performance, and formalized related open questions from a geometry perspective.

References

[1] ALEKH AGARWAL, ANIMASHREE ANANDKUMAR, PRATEEK JAIN, PRANEETH NETRAPALLI, AND RASHISH TANDON, *Learning sparsely used overcomplete dictionaries via alternating minimization*, arXiv preprint arXiv:1310.7991, (2013).

[2] ALEKH AGARWAL, ANIMASHREE ANANDKUMAR, AND PRANEETH NETRAPALLI, *Exact recovery of sparsely used overcomplete dictionaries*, arXiv preprint arXiv:1309.1952, (2013).

[3] MICHAL AHARON, MICHAEL ELAD, AND ALFRED BRUCKSTEIN, *k-svd: An algorithm for designing overcomplete dictionaries for sparse representation*, Signal Processing, IEEE Transactions on, 54 (2006), pp. 4311–4322.

[4] SAMY AIT-AOUDIA, ROLAND JEGOU, AND DOMINIQUE MICHELUCCI, *Reduction of constraint systems*, in Proceedings of Compugraphics, 1993, pp. 83–92.

[5] SANJEEV ARORA, RONG GE, AND ANKUR MOITRA, *New algorithms for learning incoherent and overcomplete dictionaries*, arXiv preprint arXiv:1308.6273, (2013).

[6] LEONARD ASIMOW AND BEN ROTH, *The rigidity of graphs*, Transactions of the American Mathematical Society, 245 (1978), pp. 279–289.

[7] TROY BAKER, MEERA SITHARAM, MENGHAN WANG, AND JOEL WILLOUGHBY, *Optimal decomposition and recombination of isostatic geometric constraint systems for designing layered materials*, Computer Aided Geometric Design, (2015).

[8] ANDERS BJÖRNER, *Subspace arrangements*, in First European Congress of Mathematics, Springer, 1994, pp. 321–370.

[9] EMMANUEL J CANDÈS, JUSTIN ROMBERG, AND TERENCE TAO, *Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information*, Information Theory, IEEE Transactions on, 52 (2006), pp. 489–509.

[10] SCOTT SHAOBING CHEN, DAVID L DONOHO, AND MICHAEL A SAUNDERS, *Atomic decomposition by basis pursuit*, SIAM journal on scientific computing, 20 (1998), pp. 33–61.

[11] DAVID L DONOHO, *Compressed sensing*, Information Theory, IEEE Transactions on, 52 (2006), pp. 1289–1306.

[12] KJERSTI ENGAN, SVEN OLE AASE, AND J HAKON HUSOY, *Method of optimal directions for frame design*, in Acoustics, Speech, and Signal Processing, 1999 IEEE International Conference on, vol. 5, IEEE, 1999, pp. 2443–2446.

[13] IOANNIS FUDOS AND CHRISTOPH M. HOFFMANN, *A graph-constructive approach to solving systems of geometric constraints*, ACM Trans. Graph., 16 (1997), pp. 179–216.

[14] MARK GORESKY AND ROBERT MACPHERSON, *Stratified morse theory*, Springer, 1988.

[15] STEVEN J GORTLER, CRAIG GOTSMAN, LIGANG LIU, AND DYLAN P THURSTON, *On affine rigidity*, Journal of Computational Geometry, 4 (2013), pp. 160–181.

[16] JACK E GRAVER, BRIGITTE SERVATIUS, AND HERMAN SERVATIUS, *Combinatorial rigidity*, vol. 2, AMS Bookstore, 1993.

[17] KIRK HALLER, AUDREY LEE-ST JOHN, MEERA SITHARAM, ILEANA STREINU, AND NEIL WHITE, *Body-and-cad geometric constraint systems*, Computational Geometry, 45 (2012), pp. 385–405.

[18] CHRISTOPH M HOFFMAN, ANDREW LOMONOSOV, AND MEERA SITHARAM, *Decomposition plans for geometric constraint problems, part ii: new algorithms*, Journal of Symbolic Computation, 31 (2001), pp. 409–427.

[19] ——, *Decomposition plans for geometric constraint systems, part i: Performance measures for cad*, Journal of Symbolic Computation, 31 (2001), pp. 367–408.

[20] CHRISTOPH M HOFFMANN AND ROBERT JOAN-ARINYO, *Symbolic constraints in constructive geometric constraint solving*, Journal of Symbolic Computation, 23 (1997), pp. 287–299.

[21] CHRISTOPH M HOFFMANN, ANDREW LOMONOSOV, AND MEERA SITHARAM, *Finding solvable subsets of constraint graphs*, in Principles and Practice of Constraint Programming-CP97, Springer, 1997, pp. 463–477.

[22] ——, *Geometric constraint decomposition*, in Geometric constraint solving and applications, Springer, 1998, pp. 170–195.

[23] BILL JACKSON AND TIBOR JORDÁN, *Pin-collinear body-and-pin frameworks and the molecular conjecture*, Discrete & Computational Geometry, 40 (2008), pp. 258–278.

[24] BILL JACKSON, TIBOR JORDÁN, AND SHIN-ICHI TANIGAWA, *Combinatorial conditions for the unique completability of low-rank matrices*, SIAM Journal on Discrete Mathematics, 28 (2014), pp. 1797–1819.

[25] DONALD J JACOBS AND BRUCE HENDRICKSON, *An algorithm for two-dimensional rigidity percolation: the pebble game*, Journal of Computational Physics, 137 (1997), pp. 346–365.

[26] CHRISTOPHE JERMANN, GILLES TROMBETTONI, BERTRAND NEVEU, AND PASCAL MATHIS, *Decomposition of geometric constraint systems: a survey*, International Journal of Computational Geometry & Applications, 16 (2006), pp. 379–414.

[27] FRANZ J KIRÁLY, LOUIS THERAN, AND RYOTA TOMIOKA, *The algebraic combinatorial approach for low-rank matrix completion*, arXiv preprint arXiv:1211.4116, (2012).

[28] AUDREY LEE AND ILEANA STREINU, *Pebble game algorithms and sparse graphs*, Discrete Mathematics, 308 (2008), pp. 1425–1437.

[29] AUDREY LEE, ILEANA STREINU, AND LOUIS THERAN, *Finding and maintaining rigid components.*, in CCCG, 2005, pp. 219–222.

[30] AUDREY LEE-ST.JOHN AND JESSICA SIDMAN, *Combinatorics and the rigidity of CAD systems*, Computer-Aided Design, 45 (2013), pp. 473 – 482. Solid and Physical Modeling 2012.

[31] PERCY S LIANG, DANIEL J HSU, AND SHAM M KAKADE, *Identifiability and unmixing of latent parse trees*, in Advances in neural information processing systems, 2012, pp. 1511–1519.

[32] ANDREW LOMONOSOV, *Graph and combinatorial analysis for geometric constraint graphs*, PhD thesis, University of Florida, 2004.

[33] JOHANNES MITTMANN, *Gröbner bases: Computational algebraic geometry and its complexity*, (2007).

[34] C ST JA NASH-WILLIAMS, *Edge-disjoint spanning trees of finite graphs*, Journal of the London Mathematical Society, 1 (1961), pp. 445–450.

[35] BRUNO A OLSHAUSEN AND DAVID J FIELD, *Sparse coding with an overcomplete basis set: A strategy employed by V1?*, Vision research, 37 (1997), pp. 3311–3325.

[36] IGNACIO RAMIREZ, PABLO SPRECHMANN, AND GUILLERMO SAPIRO, *Classification and clustering via dictionary learning with structured incoherence and shared features*, in Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on, IEEE, 2010, pp. 3501–3508.

[37] GEO REY DAVIS, *Adaptive nonlinear approximations*, PhD thesis, Courant Institute of Mathematical Sciences New York, 1994.

[38] AMIT SINGER AND MIHAI CUCURINGU, *Uniqueness of low-rank matrix completion by rigidity theory*, SIAM Journal on Matrix Analysis and Applications, 31 (2010), pp. 1621–1641.

[39] MEERA SITHARAM, *Combinatorial approaches to geometric constraint solving: Problems, progress, and directions*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 67 (2005), p. 117.

[40] DANIEL A SPIELMAN, HUAN WANG, AND JOHN WRIGHT, *Exact recovery of sparsely-used dictionaries*, in Proceedings of the Twenty-Third international joint conference on Artificial Intelligence, AAAI Press, 2013, pp. 3087–3090.

[41] PABLO SPRECHMANN AND GUILLERMO SAPIRO, *Dictionary learning and sparse coding for unsupervised clustering*, in Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on, IEEE, 2010, pp. 2042–2045.

[42] ILEANA STREINU AND LOUIS THERAN, *Sparse hypergraphs and pebble game algorithms*, European Journal of Combinatorics, 30 (2009), pp. 1944–1964.

[43] ——, *Slider-pinning rigidity: a maxwell–laman-type theorem*, Discrete and Computational Geometry, 44 (2010), pp. 812–837.

[44] WILLIAM T TUTTE, *On the problem of decomposing a graph into n connected factors*, Journal of the London Mathematical Society, 1 (1961), pp. 221–230.

[45] RENÉ VIDAL, *Subspace clustering*, Signal Processing Magazine, IEEE, 28 (2011), pp. 52–68.

[46] RENE VIDAL, YI MA, AND SHANKAR SASTRY, *Generalized principal component analysis (gpca)*, Pattern Analysis and Machine Intelligence, IEEE Transactions on, 27 (2005), pp. 1945–1959.

[47] NEIL WHITE AND WALTER WHITELEY, *The algebraic geometry of motions of bar-and-body frameworks*, SIAM Journal on Algebraic Discrete Methods, 8 (1987), pp. 1–32.

[48] NEIL L WHITE AND WALTER WHITELEY, *The algebraic geometry of stresses in frameworks*, SIAM Journal on Algebraic Discrete Methods, 4 (1983), pp. 481–511.

[49] WALTER WHITELEY, *A matroid on hypergraphs, with applications in scene analysis and geometry*, Discrete & Computational Geometry, 4 (1989), pp. 75–95.

[50] ——, *Some matroids from discrete applied geometry*, Contemporary Mathematics, 197 (1996), pp. 171–312.